TYPE DESCRIPTOR METAMODEL

Shyh-Mei F. Ho

Peter Elderon

James Rhyne

Nicholas Tindall

Tony Tsai

Cross Reference to Related Application

This application claims the benefit under Title 35, United States Code, Sections 111(b) and 119(e), relating to Provisional Patent Applications, of the filing date of United States Provisional Patent Application Serial Number 60/223,671 filed August 8, 2000 of Steven A. Brodsky and Shyh-Mei Ho for **EAI Common Application Metamodel**.

This application is also related to the following United States Patent Applications, filed on even date herewith:

COMMON APPLICATION METAMODEL by Shyh-Mei Ho, Stephen Brodsky, and James Rhyne,

COBOL METAMODEL by Shyh-Mei Ho, Nick Tindall, James Rhyne, Tony Tsai, Peter Elderon, and Shahaf Abileah.

PL/I METAMODEL by Shyh-Mei Ho, Peter Elderon, Eugene Dong and Tony Tsai.

HIGH LEVEL ASSEMBLER METAMODEL by Shyh-Mei Ho, John Ehrman, Benjamin Sheats, and Jenny Hung.

IMS TRANSACTION MESSAGES METAMODEL by Shyh-Mei Ho
and Shahaf Abileah.

IMS-MFS (MESSAGE FORMAT SERVICE) METAMODEL by
Shyh-Mei Ho, Benjamin Sheats, Elvis Halcrombe, and Chenhuei J.
Chiang.

CICS-BMS (BASIC MESSAGE SERVICE) METAMODEL by Shyh-
Mei Ho, Andy Krasun, and Benjamin Sheats.

5

10

Field of the Invention

The invention relates to exchanging instructions and/or data between applications to
signal readiness to transfer, exchange, or process data, or to establish at least one or more
parameters for transferring data between the applications, and controlling the parameters
in order to facilitate data transfer and communication. The invention further relates to
integrating dissimilar applications one executing within one platform and another
executing in another platform, e.g., multiple computers, multiple operating systems,
multiple application components, multiple development environments, multiple
deployment environments, or multiple testing and processing, establishing a dialog (e.g.,
a negotiation) with one another in order to establish connectivity for transferring data
and/or instructions between the applications so as to facilitate performing tasks on the
data or portions thereof to accomplish an overall goal. The parameters may include one
or more of format, data types, data structures, or commands.

15

20

25

Background

The growth of e-business has created a significant need to integrate legacy applications and bring them to the Internet This is because the current trend for new applications is to embrace Web standards that simplify end user application construction and scalability. Moreover, as new applications are created, it is crucial to seamlessly integrate them with

5   existing systems while facilitating the introduction of new business processes and paradigms.

Integrating new applications with existing applications is especially critical since industry analysts estimate that more than seventy percent of corporate data, including data highly

10   relevant to e-commerce, lives on mainframe computers. Moreover, while many e-commerce transactions are initiated on Windows, Mac, and Linux end user platforms, using a variety of Web browsers, and go through Windows NT and Unix servers, they are ultimately completed on mainframe computers, running mainframe applications, and impacting data stored in mainframe databases.

15

There are e-business pressures to integrate server level applications and bring them to the Internet. However, there is no complete and easy mechanism to integrate or e-business enable the applications. Integration, whether through messaging, procedure calls, or database queries, is key to solving many of today's business problems.

20

Integrating legacy applications with new software is a difficult and expensive task due, in large part, to the need to customize each connection that ties together two disparate applications. There is no single mechanism to describe how one application may allow itself to be invoked by another.

25

One consequence is an e-commerce environment of multiple applications, developed by multiple development teams, running on different platforms, with different data types, data structures, commands, and command syntax's. This environment is stitched together with application program interfaces and connectors. Connectors are an essential part of

30   the total application framework for e-commerce. Connectors match interface requirements of disparate applications and map between disparate interfaces.

3

This growing interconnection of old and new software systems and applications, has led to various middle ware applications and connector applications, interface specifications, interface definitions, and code, especially for the interconnection and interaction of

5   markup languages (such as HTML, XML, Dynamic HTML, WML, and the like), through object oriented languages such as SmallTalk and C++, with languages of legacy application server applications (such as COBOL). These interface specifications, definitions, and code should apply across languages, tools, applications, operating systems, and networks so that an end user experiences the look, feel, and responses of a

10   single, seamless application at her terminal. Instead, the proliferation of standards, protocols, specifications, definitions, and code, e.g., Common Object Request Broker (CORBA), Common Object Model (COM), Object Linking and Embedding (OLE), SOM, ORB Plus, Object Broker, Orbix, has instead created an e-commerce "Tower of Babel."

15

Examples of application integration are ubiquitous: from installing an ERP system, to updating an Operational Data Store (ODS) with IMS transactions or invoking CRM systems from MQSeries; each of these requires the same basic steps. First, a user must find the entity she wants to communicate with, then she must figure out how to invoke

20   the entity, and finally she must provide translation from one native representation to another. Today, these steps usually require manual investigation and hand coding - and leave the developers with a rat's-nest of hard-to-maintain connections between applications.

25   Attempts to remedy this situation involve application program interfaces and connectors, which are frequently built on Interface Definition Languages. Interface Definition Languages are declarative, defining application program interfaces, and, in some cases, issues such as error handling. Most Interface Definition Languages are a subset of C++, and specify a component's attributes, the parent classes that it inherits from, the

30   exceptions that it raises, the typed events that it emits, the methods its interface supports, input and output parameters, and data types. The goal of Interface Definition Languages

4

within connectors is to enable collaboration between dissimilar applications without hard coded application program interfaces.

Ideally, the interface definition language, and the connector of which it is a part, should
5   facilitate full run-time software application collaboration through such features as

- Method invocation with strong type checking,
- Run-time method invocation with greater flexibility and run time binding,
- High level language binding, with the interface separated from the
10   implementation.
- An interface repository containing real time information of server functions and parameters.

Additionally, the connector and its interface definition language, should be fast, efficient,
15   scalable, portable, support metaclasses, support syntactic level extensions, and support semantic level extensions.

Summary of the Invention

20   The problems associated with integrating new applications, for example, e-commerce applications, with legacy applications are obviated by the Common Application Metamodel tool, method, and system described herein. The Common Application Metamodel method, tool, and system of the invention facilitate tooling solutions, data translation, and communication and collaboration between dissimilar and disparate
25   applications, as well as full run-time software application collaboration through an interface with the application server interface domain. This is accomplished through metadata interchange information, method invocation with strong type checking, run-time method invocation, run time binding, and high level language binding, with the interface separated from the implementation, and an interface repository containing real
30   time information of client and server interface parameters.

SVL920010038US1

Additionally, the tool, method, and system of the invention provide fast, efficient, and scalable interconnectivity independently of any tool or middleware, are reusable and portable, and support metaclasses, syntactic level extensions, and semantic level extensions, and are independent of any particular tool or middleware.

5

The Common Application Metamodel tool, method, and system is especially useful for providing a data transformer that is bi-directional between a client application and a server application, transmitting commands and data both ways between, for example, a Java, HTML, XML, C, or C++ application and a COBOL, PL/I, or High Level

10   Assembler application, or, between an HTML or XML application and a Java, C, or C++ application, or between a Java application and a C or C++ application.

One embodiment of the invention is a method of processing a transaction on or between an end user application and one or more application servers. The method comprises the

15   steps of initiating the transaction on the end user application in a first language with a first application program, transmitting the transaction to the server, and converting the transaction from the first language of the first end user application to a language running on the application server. Typically, as described above, the client will be a thin client or a Web browser, the application running on the client will be a Web browser application

20   or a thin client connectivity application, and the language of the client application will be Java, C, C++, or a markup language, as HTML or a derivative of HTML, such as XML or Dynamic HTML or WML, or the like, and the language running on the server may be COBOL, PL/I, HLASM (High Level Assembler) or the like. The invention facilitates transformers which convert the transaction from the first language of the end user

25   application to a language running on the application server. After conversion, the converted transaction is processed on the application server.

The application processes the request and then sends the response from the application server back to the end user application. Typically, as described above, the application

30   server will be running a COBOL based application, and the client will be a thin client written in Java or C or C++, or a Web browser, running a Web browser application or a

6

thin client connectivity application, in a markup language, as HTML or a derivative of HTML, such as XML or Dynamic HTML, or the like. The invention provides data transformers which convert the response from the language or languages running on the application server or servers to the first language of the first end user application.

5

The end user application and the application server have at least one data transformer between them. In this way, the steps of (i) converting the request from the first language of the first end user application as a source language to the language running on an application server as a target language, and (ii) converting the response from the language
10   running on the application server, as a subsequent source language, back to the first language of the first end user application, as a subsequent target language, each comprise the steps of invoking type descriptor and language metamodels of respective source and target languages, populating the metamodels with each of the respective source and target languages' data items and types, and converting the source language to the target
15   language.

The end user application is, frequently, a web browser or a thin client. When the end user application is a Web browser, the end user is connected to the application server through a web server. According to a further embodiment of the invention, the web server may
20   comprise the connector, or data transformer. The data transformer integrated with the Web server may directly convert the request, transaction, or message from a browser oriented form to an application server language or to an intermediate, business or commerce oriented markup language, such as XML.

25   The CAM metamodel used to construct the converter comprises an invocation metamodel, an application domain interface metamodel, a language metamodel, and a type descriptor metamodel. Exemplary invocation metamodel includes information chosen from the group consisting of message control information, security data, transactional semantics, trace and debug information, pre-condition and post-condition
30   resources, and user data, etc. Exemplary application domain interface metamodel comprises information chosen from input parameter signatures, output parameter

7

signatures, and return types. Application domain interface metamodel uses one or more language metamodels, such as COBOL and PL/I metamodels.

5 The type descriptor metamodel defines physical realizations, storage mapping, data types, data structures, and realization constraints.

10 The method of the invention is applicable to situations where one of the source or target languages is object oriented, and the other of the target or source languages is not object oriented. In this situation, the language metamodel and the type descriptor metamodel together map encapsulated objects of the object oriented language into code and data of the language that is not object oriented. Additionally, the language metamodel and the type descriptor metamodel maps object inheritances of the object oriented language into references and pointers in the language that is not object oriented. The method of the invention is also applicable to situations where different object oriented languages are 15 running on different platforms, and encapsulated objects of the source language (code and data) are mapped into encapsulated objects of the target language. The method of the invention is also applicable where different procedural languages are running on different platforms or applications and commands and data of the source procedural language are mapped into the target procedural language.

20 According to the method of the invention, there may be a plurality of applications for vertical (sequential, conditional, or dependent) processing, for horizontal (parallel in time) processing, or both horizontal and vertical processing. This is to support rich transactions to and through multiple hierarchical levels and multiple parallel sequences of 25 processing. This may be the case in business to business transactions drawing upon financial, manufacturing, scheduling, supply, and shipping databases and servers, and utilizing various commercial security instruments.

A further aspect of the invention is a client-server processing system having a client, a 30 server, and at least one transformer between the client and one or more servers,

SVL920010038US1

A still further aspect of the invention is a processing system configured and controlled to interact with a client application. In this aspect of the invention, the system comprises, a server, and at least one transformer between the server and the client application, where the client has an end user application, and is controlled and configured to initiate a

5   request with the server in a first language with a first application program and to transmit the request through a transformer to the server or servers. The server processes the request in a second software application, using a second language, and returns a response to the client through a transformer.

10   A further aspect of the invention is a groupware system having a plurality of e-mail enabled end user applications, such as e-mail, word processing, spreadsheet, simple database management (such as Lotus Approach or Microsoft Access), graphics and graphics editing, audio and audio editing, and computer-telephony integration ("CTI"), along with client level content database client services and content replication client

15   services. Groupware integrates these e-mail enabled applications through one or more transformers and application program interfaces with transport services, directory services, and storage services, including content servers and replication servers. The groupware system is configured and controlled to communicate among disparate end user applications, among disparate servers, and between disparate servers and end user

20   applications. The groupware system comprises at least one transformer between a server and an end user application. The end user application is controlled and configured to participate with a server in a first language of a first application program and the server is configured and controlled to participate with the client in a second language of a second program.

25

The transformer is configured and controlled to receive a request from the end user application, and convert the request from the first language of the first end user application to a language running on the server. The server is configured and controlled to receive the converted request from the transformer and process the request in a second

30   language with a second application program residing on the server, and to thereafter transmit a response through a transformer back to the end user application.

9

A still further embodiment of the invention is the provision of rich transaction processing. Rich transactions are nested transactions that span to, through, and/or across multiple servers. The spanning across nested servers may be horizontal, that is parallel dependent

5   transactions, or vertical, that is, serial dependent transactions. Rich transactions may be long lived, on-going transactions, or complex business-to-business transactions, especially those with multiple dependencies or contingencies, volume and prompt payment discounts, late delivery and late payment penalties, and with financial processing, such as electronic letters of credit, electronic bills of lading, electronic

10   payment guarantees, electronic payment, escrow, security interests in the goods, and the like. In a rich transaction environment, some transaction servers may be positioned as clients with respect to other transactions for certain sub transactions making up the rich transaction.

15   A still further embodiment of the invention is a tool, that is, a software developer's kit, characterized in that the program product is a storage medium (as a tape, floppy disks, a CD-ROM, or a hard drive or hard drives on one of more computers) having invocation metamodels, application domain interface metamodels, and language metamodels, and computer instructions for building a metamodel repository of source and target language

20   metamodels. The program product also contains computer instructions for building connector stubs from the metamodels. The program product further carries computer instructions to build a transformer.

While the invention has been described in summary form as having a single level of

25   connectors, it is, of course, to be understood that such connectors may be present at various levels in the processing hierarchy, for example between Web Clients and Web servers, between web servers and application servers, between application servers and database servers, and between application servers or database servers or both and various specialized repositories.

30

SVL920010038US1

It is also to be understood, that while the invention has been summarized in terms of individual clients and individual servers, there may be multiple clients, multiple servers, and applications that function as both clients and servers, as exemplified by groupware applications, and there might be multiple parallel lines and/or multiple hierarchical levels

5 of application servers, data servers, and databases, as in systems for rich transactions.

The Figures

Various elements of the invention are illustrated in the FIGURES appended hereto.

10

FIGURE 1 illustrates a system with multiple application components, including a Netscape Internet Explorer browser, Net.Commerce on a Sun Solaris server, Oracle and DB2 on a database server, SAP running on AIX, a CICS 390 server, an IMS 390 server, DB2 and DL/I on a S/390 platform, a Windows 200 client, and Baan running on an HP

15 Unix server.

FIGURE 2 illustrates the roles of message sets, SQL stored procedures, legacy applications, and programming languages as inputs to the metadata repository of the Common Application Metamodel to facilitate enterprise application integration at run

20 time.

FIGURE 3 illustrates that the Common Application Metamodel of the invention consists of three kinds of metamodels, i.e., an invocation metamodel, an application-domain interface metamodel, and a language metamodel. For any given application-domain

25 metamodel it may use one or many language metamodels, and there could be zero or many invocation metamodels.

FIGURE 4 illustrates an IMS OTMA metamodel, with an OTMA Invocation Metamodel, an IMS Transaction Message Metamodel application interface, which could use a

30 COBOL Metamodel, a C Metamodel, or other language metamodels.

11

FIGURE 5 illustrates how a tool can be used to generate an XML document describing application program interface. First, an object model, i.e., a CAM metamodel, is created to capture interface definitions about an application server. Then a tool reads and parses the source definitions of an application program and generates an XML document by

5   retrieving the object model's information from a repository.

FIGURE 6 illustrates a development phase scenario where a Common Application Metamodel Rose file, e.g., a COBOL metamodel, a PL/I metamodel, an MFS metamodel, a BMS model, or the like is read into a toolkit, to generate a DTD and XML schema and

10   Java code for a Rose model. A source file of an application, as a COBOL source file, a PL/I source file, an MFS source file, a BMS source file, or the like, is read into an importer. The importer parses the source code and generates, as output, an XMI instance file, i.e., XML documents, by reading in the Java code of the Rose model of the application source files.

15

FIGURE 7 illustrates a metamodel for application interfaces, which enables integration of application components into an event based messaging model, including flow models. The flow and messaging middle invokes applications through the application interface. These interfaces are access points to the applications through which all input and output

20   is connected to the middleware. The interfaces are described in terms of the Application Interface Metamodels. Transformation processing according to he metamodel could take place in source/client applications, target applications, or a gateway.

FIGURE 8 illustrates the application of the Common Application Metamodel during

25   execution time. As shown, the CAM model facilitates connectivity between a back-end IMS application and a Web file (e.g., SOAP complaint XML documents). This is accomplished by using information captured in the model to perform data transformations from one platform to another in a mixed language environment shown.

30   FIGURE 9 illustrates TDLang base classes, where the Type Descriptor metamodel is used as a recipe or template for runtime data transformation, with the language specific

SVL920010038US1

metamodel for overall data structures and field names, and without duplicating the aggregation associations present in the language model.

FIGURE 10 illustrates an MOF class instance at the M2 level as a Type Descriptor Metamodel.

FIGURE 11 illustrates the association between a TDLangElement of the TDLang Metamodel with a Platform Compiler Type of the Type Descriptor Metamodel.

FIGURE 12 illustrates various enumeration types for the Type Descriptor Metamodel, e.g., sign coding, length encoding, floating type, accessor, packed decimal sign, and bitModelPad.

FIGURE 13 illustrates a simplified network configuration for a "rich" transaction where, for example, an order is entered at a terminal, and is passed to and through a Web server to a manufacturer's application server. The manufacturer's application server searches through it's own database server, as well as its vendors' dissimilar and incompatible database servers and databases, transparently connected by the connectors described herein, querying for statuses, prices, and delivery dates, of components, and placing orders for needed components to satisfy the order.

FIGURE 14 illustrates a group ware session spread across multiple group ware applications running on multiple, disparate platforms, and connected using the common application metamodel described herein.

FIGURE 15 illustrates a commercial transaction where real goods are shipped from a seller to a buyer, and various forms of electronic payment and secured electronic payment are used by the buyer to pay the seller, with banks and financial institutions connected using the common application metamodel described herein.

Detailed Description of the Invention

13

Definitions. As used herein the following terms have the indicated meanings.
"Handshaking" is the exchange of information between two applications and the resulting
agreement about which languages, capabilities, and protocols to use that precedes each

5   connection.

An "application program interface" (API) is a passive specific method prescribed by a
computer operating system or by another application program by which a programmer
writing an application program can make requests of the operating system or another

10   application. Exemplary is SAX (Simple API for XML), an connector that allows a
programmer to interpret a Web file that uses the Extensible Markup Language, that is, a
Web file that describes a collection of data. SAX is an *event-driven* interface. The
programmer specifies an event that may happen and, if it does, SAX gets control and
handles the situation. SAX works directly with an XML parser.

15

A "connector" as used herein is a dynamic, run-time, interface between platforms that
stores the functions and parameters of the target platform or program, and binds with the
target platform program in real time.

20   A "stub" is a small program routine that provides static interfaces to servers. Precompiled
stubs define how clients invoke corresponding services on the server. The stub
substitutes for a longer program on the server, and acts as a local call or a local proxy for
the server object. The stub accepts the request and then forwards it (through another
program) to the remote procedure. When that procedure has completed its service, it

25   returns the results or other status to the stub which passes it back to the program that
made the request. Server services are defined in the stub using an Interface Definition
Language ("IDL"). The client has an IDL stub for each server interface that it accesses
and includes code to perform marshaling. Server stubs provide static interfaces to each
service exported by the server.

30

14

"CICS" (Customer Information Control System) is the online transaction processing program from IBM that, together with the Common Business Oriented Language programming language, is a set of tools for building customer transaction applications in the world of large enterprise mainframe computing. Using the programming interface

5   provided by CICS to write to customer and other records (orders, inventory figures, customer data, and so forth) in a CICS, a programmer can write programs that communicate with online users and read from a database (usually referred to as "data sets") using CICS facilities rather than IBM's access methods directly. CICS ensures that transactions are completed and, if not, it can undo partly completed transactions so that

10   the integrity of data records is maintained. CICS products are provided for OS/390, UNIX, and Intel PC operating systems. CICS also allows end users to use IBM's Transaction Server to handle e-business transactions from Internet users and forward these to a mainframe server that accesses an existing CICS order and inventory database.

15   "IMS" (Information Management System) is the system from IBM that, together with IBM's Enterprise Systems Architecture (IMS/ESA) provides a transaction manager and a hierarchical database server.

"MQ" is the MQSeries IBM software family whose components are used to tie together

20   other software applications so that they can work together. This type of application is often known as business integration software or middleware. Functionally, MQSeries provides a communication mechanism between applications on different platforms, an integrator which centralizes and applies business operations rules, and a workflow manager which enables the capture, visualization, and automation of business processes.

25   MQSeries connects different computer systems, at diverse geographical locations, using dissimilar IT infrastructures, so that a seamless operation can be run. IBM's MQSeries supplies communications between applications, and between users and a set of applications on dissimilar systems. Additionally, MQSeries' messaging scheme requires the application that receives a message to confirm receipt. If no confirmation

30   materializes, the message is resent by the MQSeries.

15

"Rose" is an object-oriented Unified Modeling Language (UML) software design tool intended for visual modeling and component construction of enterprise-level software applications. It enables a software designer to visually create (model) the framework for an application by blocking out classes with actors (stick figures), use case elements

5 (ovals), objects (rectangles) and messages/relationships (arrows) in a sequence diagram using drag-and-drop symbols. Rose documents the diagram as it is being constructed and then generates code in the designer's choice of C++, Visual Basic, Java, Oracle8, Corba or Data Definition Language.

10 **Common Application Metamodel Overview**. The Common Application Metamodel (CAM) brings interconnectivity to the environment illustrated in FIGURE 1. FIGURE 1 illustrates a typical system 101 with multiple application components, including a Netscape Internet Explorer browser 103, Net.Commerce 105 on a Sun Solaris server 107, Oracle 109 and DB2 111 on a database server 113, SAP 115 running on AIX 117, a CICS

15 390 server 119, an IMS 390 server 121, DB2 123 and DL/I 125 on a S/390 platform 127, a Windows 2000 client 129, and Baan 131 running on an HP Unix server 133. The Common Application Metamodel (CAM) is metadata interchange method, tool, and system for marshaling and applying information needed for accessing enterprise applications, such as in FIGURE 1, in a source language and converting them to a target

20 language. CAM consists of language metamodels and application domain interface metamodels, as shown in FIGURE 2, which illustrates the roles of message sets 203, SQL stored procedures 205, legacy applications 207, and programming languages 209 as inputs to the metadata repository 211 of the Common Application Metamodel to facilitate enterprise application integration 221.

25

Exemplary metamodels include C, C++, Java, COBOL, PL/I, HL Assembler, IMS transaction messages, IMS MFS, CICS BMS, and MQSeries messages models, as shown in FIGURE 3, which illustrates the Common Application Metamodel of the invention, with an invocation metamodel 301, an application-domain interface metamodel 303, and

30 a language metamodel 305.

SVL920010038US1

FIGURE 4 illustrates an IMS OTMA application interface metamodel 411, with an OTMA Invocation Metamodel 421, an IMS Transaction Message Metamodel 423, a COBOL Metamodel 425, and a C Metamodel 427.

5   FIGURE 5 illustrates the flow of information from an existing application 501, through an interface 503 to an object model containing application interface metadata. This application interface metamodel is stored in the metadata repository 505, and, at an appropriate time, retrieved from the metadata repository 505, combined with a source program 507 in a generation tool 509, and used to generate a target file 511, as an XML

10   file, i.e., an XMI instance file. CAM is highly reusable and independent of any particular tool or middleware.

Development Stage. With CAM, tooling can now easily provide solutions to access enterprise applications, e.g. IMS applications. By parsing each source file and generating

15   XML documents based on the CAM model, COBOL copybook, PL/I copybook, MFS Source, BMS Source, etc., tools can provide connector solutions to IMS, and CICS, etc.

In this regard, FIGURE 6 illustrates a development phase scenario where a Common Application Metamodel Rose file 601, e.g., a COBOL metamodel, a PL/I metamodel, an

20   MFS metamodel, a BMS model, or the like is read into a toolkit 603, to generate a DTD and schema for a Rose model and Java code for a Rose model 605. A source file of an application 607, as a COBOL source file, a PL/I source file, an MFS source file, a BMS source file, or the like, and the Java code for the Rose model 609 are read into an Importer 611. The importer parses the source code and provides, as output, an XMI

25   instance file 613, i.e., XML documents, of the application source files.

FIGURE 7 shows a CAM metamodel for application interfaces. This Figure depicts a runtime connector 701 with invocation and transformation capabilities, interfacing with an existing application program 703 through an interface 705 containing the existing

30   application program's interface definition, in accordance with the application interface

17

metamodel 707. The Application Interface metadata is stored in a metadata repository 709.

The flow and messaging middleware 713 invokes applications 703 through the
5 application interfaces 705. These interfaces 705 are the access points to the applications 703 through which all input and output is connected to the middleware 713. The interfaces 705 are described in terms of the Application Interface Metamodel. Transformation processing according to the metamodel could take place in source/client applications, target applications, or a gateway.
10

Because CAM also provides physical representation of data types and storage mapping to support data transformation in an enterprise application integration environment, it enables Web services for enterprise applications.

At development time CAM captures information that facilitates:
15   a). connector and/or connector-builder tools,

b). data driven impact analysis for application productivity and quality assurance, and

c). viewing of programming language data declarations by developers.

The CAM metamodel files are inputs to toolkits used to generate DTD files, XML
20 schemas, and Java classes which represent the CAM model. Importers parse each source file (e.g. COBOL or PL/I copybook, MFS source, and BMS, etc.), and then generate XML documents (i.e. XML instance files) based on Java classes generated by the XMI/MOF2 toolkit.

25 Run Time. At run time CAM provides information which facilitates transformation in an enterprise application integration environment where it provides data type mapping between mixed languages, facilitates data translations from one language and platform domain into another.

30 FIGURE 8 illustrates the application of the Common Application Metamodel during run time. As shown, SOAP compliant XML documents 803 are received in, for example,

18

IBM WebSphere middleware, 805, which contains an IMSConnector for Java 807, and is in contact with an XML Repository 809, containing the XMI instance files for the CAM model. The IBM WebSphere middleware sends the transformed file to the IMS system 811, which contains an instance of IMS Connect 813 and the IMS transactional

5   application program 815. CAM facilitates connectivity between the back-end IMS application 815 and the Web file (e.g., SOAP compliant XML documents) 803. The CAM accomplishes this by using CAM model information (from repository 809) to perform data transformations from one platform to another in the mixed language environment shown.

10

Type Descriptor Metamodel. One important feature provided by CAM is the Type Descriptor metamodel. The Type Descriptor metamodel defines the physical realization, storage mapping, and the constraints on the realization (such as justification). This metamodel provides a physical representation of individual fields of a given data

15   structure. When supporting data transformation in an enterprise application integration environment, the model provides data type mapping between mixed languages. It also facilitates data translations from one language and platform domain into another. The metamodel is used for runtime data transformation (or marshaling) with a language-specific metamodel for overall data structures and field names.

20

1.   Common Application Metamodel for Application Interfaces

The interconnection of disparate and dissimilar applications running on different software platforms, as shown in FIGURE 1, with different operating systems, physical platforms,

25   and physical realizations is accomplished through connectors that incorporate the interconnection metadata. Connectors are a central part of the application framework for e-business. The end user demand is to connect to anything interesting as quickly, and as easily, as possible.

30   A connector is required to match the interface requirements of the adapter and the legacy application. It is also required to map between the two interfaces. Standardized

metamodels for application interfaces presented herein allow reuse of information in multiple connector tools. These standardized metamodels not only reduce work to create a connector, but also reduce work needed to develop connector builder tools.

5   The connectors built using the common application metamodel of our invention provide interoperability with existing applications. The connectors support leveraging and reuse of data and business logic held within existing application systems. The job of a connector is to connect from one application system server "interface" to another. Therefore, an application-domain interface metamodel describes signatures for input and

10   output parameters and return types for a given application system domain (e.g. IMS, MQSeries); it is not for a particular IMS or MQSeries application program. The metamodel contains both syntactic and semantic interface metadata.

1.   a.   End-to-End Connector Usage Using Common Application Metamodel

15

The Common Application Metamodel (CAM) consists of meta-definitions of message signatures, independent of any particular tool or middleware. Different connector builder tools can use this information to ensure the "handshaking" between these application programs, across different tools, languages, and middleware. For example, if you have to

20   invoke a MQSeries application, you would need to build a MQ message using data from a GUI tool and deliver it using the MQ API. Similarly, when you receive a message from the MQSeries application, you would need to get the buffer from MQSeries, parse it and then put it into a GUI tool data structure. These functions can be designed and implemented efficiently by a connector builder tool using CAM as standardized

25   metamodels for application interfaces.

CAM can be populated from many sources, including copy books, to generate HTML forms and JavaServer Page (JSP) for gathering inputs and returning outputs. An example of a connector as depicted in the previous figure is that the flow and message middleware

30   makes a function call to an enterprise application by calling the connector which then calls the enterprise application API. The connector does language and data type

SVL920010038US1

mappings, for example, to translate between XML documents and COBOL input and output data structures based on CAM. Connectors and CAM provide the end-to-end integration between the middleware and the enterprise applications.

5   Using IMS as an example. Let's say that you must pass an account number to an IMS transaction application program from your desktop to withdraw $50.00. With CAM and a connector builder tool, you will first generate an input HTML form and an output JSP; and develop a middleware code necessary to support the request. The desktop application fills the request data structure (i.e. an input HTML form) with values and calls the

10   middleware. The middleware service code will take the data from the GUI tool, build an IMS Connect XML-formatted message, and deliver the message to the IMS gateway (i.e. IMS Connect) via TCP/IP. IMS Connect translates between the XML documents and the IMS message data structures in COBOL using the metadata definitions captured in CAM. It then in turn sends the IMS message data structures to IMS via Open Transaction

15   Manager Access (OTMA). The IMS COBOL application program runs, and returns the output message back to the middleware service code via IMS Connect. The middleware service code gets the message and populates the output JSP page (i.e. previously generated GUI tool reply data structures) with the reply data. The transaction output data will then be presented to the user.

20

2.   Common Application Metamodel

CAM is used to describe information needed to easily integrate applications developed in common programming models with other systems. The CAM metamodel can be used

25   for both synchronous and asynchronous invocations.

2.   a.   Common Application Metamodel

The common application metamodel depicted as follows consists of an invocation

30   metamodel and an application-domain interface metamodel which uses language

SVL920010038US1

metamodels. For any given application-domain interface metamodel, it may use one or many language metamodels, but, there could be zero or more invocation metamodels.

The common connector metamodel is illustrated in FIGURE 3. It has an Invocation
5   Metamodel 301, an Application-Domain Interface Metamodel 303, and a Language Metamodel 305.

2.   a.   i.   Invocation Metamodel

10   The invocation metamodel 301 consists of one or more of the following possible metadata elements. However, for a particular invocation, it could include only one or many of the following metadata elements.

Message-control information. This includes message control information, such as the
15   message connection name, message type, sequence numbers (if any), and various flags and indicators for response, commit-confirmation, and processing options by which a client or server can control message processing to be synchronous or asynchronous, etc.

The connection name can be used by the application system server to associate all input
20   and output with a particular client. The message type specifies that the message is a response message; or that commit is complete. It can also indicate server output data to the client, or client input data to the server.

Security data -- This includes authentication mechanism, and security data, e.g. digital
25   certificates, identity, user name and password, etc. It may also include authorization information to indicate whether the data can be authorized via a role based or ACL (access control list) based authorization.

Transactional semantics -- This will carry transaction information, e.g. local vs. global
30   transaction; two-phase commit vs. one-phase commit, and transaction context, etc.

SVL920010038US1

Trace and debug -- Trace and debugging information are specified as part of the metamodel.

Precondition and post-condition resource -- This describes application state precondition

5   and post-condition relationships.

User data -- This includes any special information required by the client. It can contain any data.

10   **2.   a.   ii.   Application-Domain Interface Metamodel**

The application-domain interface metamodel 303, as discussed earlier, describes signatures for input and output parameters and return types for application system domains.

15

2.   a,   iii.   Language Metamodel

The language metamodel 305, e.g. COBOL metamodel, is used by enterprise application programs to define data structures (semantics) which represent connector interfaces. It is

20   important to connector tools to show a connector developer the source language, the target language, and the mapping between the two. The CAM language metamodel also includes the declaration text in the model which is not editable (i.e. read-only model). Because the connector/adapter developer would probably prefer to see the entire COBOL data declaration, including comments and any other documentation that would help

25   him/her understand the business role played by each field in the declaration.

The language metamodel is also to support data driven impact analysis for application productivity and quality assurance. (But, it is not the intention of the CAM to support reproduction of copybooks.)

30

The language metamodels describing connector data are listed as follows:

SVL920010038US1

* C

* C++

* COBOL

5 * PL/I

2.   a.   iv.   Type Descriptor Metamodel

The Type Descriptor metamodel is language neutral and defines the physical realization,

10   storage mapping and the constraints on the realization such as justification. This

metamodel provides physical representation of individual fields of a given data structure.

The type descriptor metamodel is to support data transformation in an enterprise

application integration environment to provide data types mapping between mix

languages. It also facilitates data translations from one language and platform domain

15   into another. This metamodel will be used as a recipe for runtime data transformation (or

marshaling) with language specific metamodel for overall data structures and fields

names.

3.   An Example of Common Connector Metamodel

20

IMS OTMA (Open Transaction Manager Access) is a transaction-based, connectionless

client/server protocol within an OS/390 sysplex environment. An IMS OTMA transaction

message consists of an OTMA prefix, plus message segments for input and output

requests. Both input and output message segments contain llzz (i.e. length of the segment

25   and reserved field), and application data. Only the very first input message segment will

contain transaction code in front of the application data. IMS transaction application

programs can be written in a variety of languages, e.g. COBOL, PL/I, C, and Java, etc.

Therefore, the application data can be in any one of these languages.

30   As shown in FIGURE 4, an IMS OTMA connector metamodel 401 is composed of an

invocation metamodel 403 and an IMS transaction message metamodel 405, as well as a

24

COBOL metamodel 407 and a C metamodel 409. As depicted in FIGURE 4, the invocation metamodel 401 is the OTMA prefix, and the IMS transaction message metamodel 405 is the application-domain interface metamodel for the IMS application system which uses language metamodels. Metamodels for COBOL 407 and C 409 are

5   shown.

4.   Type Descriptor Metamodel

The type descriptor metamodel presents a language and platform independent way of

10   describing implementation types, including arrays and structured types. This information is needed for marshaling and for connectors, which have to transform data from one language and platform domain into another. Inspections of the type model for different languages can determine the conformance possibilities for the language types. For example, a long type in Java is often identical to a binary type (computational-5) in

15   COBOL, and if so, the types may be inter-converted without side effect. On the other hand, an alphanumeric type in COBOL is fixed in size and if mapped to a Java type, loses this property. When converted back from Java to COBOL, the COBOL truncation rules may not apply, resulting in computation anomalies. In addition, tools that mix languages in a server environment (e.g., Java and COBOL in CICS and IMS) should find it useful

20   as a way to determine how faithfully one language can represent the types of another.

Therefore, an instance of the type descriptor metamodel describes the physical representation of a specific data type for a particular platform and compiler.

25   ## 4.   a.   TDLang Metamodel

The TDLang metamodel serves as base classes to CAM language metamodels by providing a layer of abstraction between the Type Descriptor metamodel and any CAM language metamodel. All TDLang classes are abstract and common to all the CAM

30   language metamodels. All associations between TDLang classes are marked as "volatile," "transient," or "derived" to reflect that the association is derived from the language

25

metamodel. The TDLang model does not provide any function on its own, but it is the type target for the association from the Type Descriptor metamodel to the language metamodels.

5   FIGURE 9 illustrates the structure of the TDLang Metamodel, with the TDLangClassifier 501, the TDLangComposedType 503 and the TDLangElement 505.

With the TDLang base classes, the Type Descriptor metamodel can be used as a recipe for runtime data transformation (or marshaling) with the language-specific metamodel for
10   overall data structures and field names, without duplicating the aggregation associations present in the language model.

4.   b.   Type Descriptor Metamodel

15   This metamodel is a MOF Class instance at the M2 level. FIGURE 10 shows the relationships within the type descriptor Metamodel, including the PlatformCompilerType 601, the InstanceTDBase 603, the ArrayTD 605, the AggregateInstanceTD 607, the Simple InstanceTD 609, and the InstanceType 611. The InstanceType 611 comprises definitions of the StringTD 613, the AddressTD 615, the NumberTD 617, and the
20   FloatTD 619. FIGURE 11 illustrates a higher level view of the TDLanguageElement and the PlatformCompilerType 601. FIGURE 12 illustrates enumerations of signCoding 801, lengthEncoding 803, floatType 805, accessor 807, packedDecimalSign 809, and bitModePad 811.

25   **4.   c.   Type Descriptor and Language models**

The Type Descriptor model is attached to the CAM Language model by a navigable association between TDLangElement and InstanceTDBase. TDLangElement is the base language model type used to represent a declared data item, i.e., an instance of a type.
30   InstanceTDBase is the base Type Descriptor model type used to represent the

SVL920010038US1

implementation-specific instance of this same declared data item. InstanceTDBase is abstract; only one of its subtypes may be instantiated.

It is possible that a data item declared in a programming language may have different

5   implementations. These differences are induced by hardware platform, system platform, and compiler differences. This possibility is modeled by the PlatformCompilerType model type. The association between TDLangElement and PlatformCompilerType is many to one, and the association between PlatformCompilerType and InstanceTDBase is one to one. To navigate from the language model, it is necessary to know what

10   PlatformCompilerType is to be assumed. It is possible that an implementation, upon importing a model instance, will wish to remove from the model the PlatformCompilerType instances that are not of interest.

The association between TDLangElement and InstanceTDBase is modeled in this manner

15   to allow for extending the model to include an association between PlatformCompilerType and a new type that more fully describes the hardware platform, the system platform, and the compiler.

Data element instances may be defined as repeating groups or arrays. This is modeled as

20   a one to many association between InstanceTDBase and the ArrayTD model type. There would be one ArrayTD instance in this association for each dimension, subscript, or independent index of the data element. These instances hold information about the bounds and accessing computations.

The association is ordered in the same order as the corresponding association in the

25   language model, and reflects the syntactic ordering of the indices as defined by the programming language. The rationale for this choice is the resulting equivalence of navigation and processing algorithms between the language model and the Type Descriptor model. Another choice, perhaps more advantageous to marshaling engines, would be to have the ordering of the indices from the smallest stride to the largest. This

30   allows a marshaling engine to process the array in its natural storage order, assuming it is

SVL920010038US1

laid out in the usual contiguous fashion. A marshaling engine can compute this order by re-sorting the association targets according to the stride formulas if desired.

Array information may be a complex property of the data element or of its type, and

5   various languages and programming practices seem to fall on either side. The typedef facility of C and C++ allows the definition of some array types from typedefs, but only where the array definitions are applied to the topmost elements of typedef aggregates. For example, consider the following typedef:

10
```
typedef struct {
        int A;
        struct {
                int C;
                char D;
15              struct {
                        int F;
                        int G;
                } E;
        } B;
20 } X;
```

This typedef can be used to create a new typedef for a fixed size array, e.g.

typedef X Q[10];

25

But it is not possible to create a new typedef from X that makes any of the subcomponents of X, e.g., D or E, into an array. This example and many others point out the unclear status of array definitions in typed languages.

30   An InstanceTDBase type has two concrete subtypes, SimpleInstanceTD and AggregateInstanceTD. SimpleInstanceTD models data elements without subcomponents, while AggregateInstanceTD models data elements with subcomponents. To find the subcomponents of an AggregateInstanceTD, one must navigate back to the corresponding data element declaration in the CAM language model. There, the association between an

35   aggregate type and its subcomponents may be navigated, leading to a set of

28

subcomponent data elements, each of which has one or more corresponding instances in the Type Descriptor model. This introduces some model navigation complexity, but avoids duplicating the aggregation hierarchy in both the language and the Type Descriptor models. The additional processing complexity of traversal is not great, and

5   considerable simplification is obtained in algorithms that would modify the model to add, delete or rearrange subcomponents in an aggregation.

A SimpleInstanceTD model type is also associated one to one with a BaseTD model type. The BaseTD model type is specialized to hold implementation information that is

10   common for all data elements of the same language type. The information that describes a 32-bit signed binary integer on a specific hardware/software platform is thus instantiated only once in a given model instantiation, no matter how many data elements may be declared with this type.

15   One may contemplate an association between TDLangClassifier and BaseTD matching the association between TDLangElement and InstanceTDBase. However, this is problematic in that constructions that the language regards as simple types (e.g., strings) may not map directly to simple hardware/software types. Rather than introduce more mechanisms into the Type Descriptor model to describe string implementations, a

20   specialization of BaseTD is utilized which describes the common string implementations. Various attributes in the TypeDescriptor model are suffixed with the string "formula." These attributes contain information that may in some cases be impossible to compute without access to data created only at run-time. An example is the current upper bound of a variable-sized array or the offset to an element that follows another element whose size

25   is only known at run-time. Such information could be included as values in a model instance, but this would require a model instance for each run-time instance, and would mean that the model could only be constructed at run-time, requiring the model definition to include factories and other apparatus to create model instances at run-time. A model that can be constructed from platform and compiler knowledge is much more useful, and

30   the formulas provide a way to define concrete values when the run-time information is available. These formulas may be interpreted by marshaling engines, or they may be used

to generate marshaling code, which is loaded and executed by the marshaling engine on demand.

4. d. Formulas

5

As used in connection with formulas, "field" refers to a component of a language data structure described by the Type Descriptor model, while "attribute" denotes part of the model, and has a value representing a "property" of the field. Thus the value of a field means a run-time value in a particular instance of a language data structure, whereas the

10 value of an attribute is part of the description of a field in a language data structure, applies to all instances of the data structure, and is determined when the data structure is modeled.

For most attributes in an instance of the Type Descriptor model, the value of the attribute

15 is known when the instance is built, because the properties of the fields being described, such as size and offset within the data structure, are invariant. But if a field in a data structure is defined using the COBOL OCCURS DEPENDING ON construct or the PL/I Refer construct, then some properties of the field (and properties of other fields that depend on that field's value) cannot be determined when the model instance is built.

20

Properties that can be defined using these language constructs are string lengths and array bounds. A property that could indirectly depend on these language constructs is the offset of a field within a structure, if the field follows a variable-size field.

25 In order to handle these language constructs, properties of a field that could depend on these constructs (and thus the values of the corresponding attributes), are defined with strings that specify a formula that can be evaluated when the model is used.

However, if a property of a field is known when the model instance is built, then the

30 attribute formula simply specifies an integer value. For example, if a string has length 17, then the formula for its length is "17".

The formulas mentioned above are limited to the following:

* Unsigned integers

* The following arithmetic integer functions

5
```
neg(x)   := -x       // prefix negate
add(x,y) := x+y      // infix add
sub(x,y) := x-y      // infix subtract
mpy(x,y) := x*y      // infix multiply
div(x,y) := x/y      // infix divide
```
10
```
max(x,y) := max(x,y)
min(x,y) := min(x,y)
mod(x,y) := x mod y
```

The mod function is defined as mod(x,y) = r where r is the smallest non-negative integer

15 such that x-r is evenly divisible by y. So mod(7,4) is 3, but mod(-7,4) is 1. If y is a power

of 2, then mod(x,y) is equal to the bitwise-and of x and y-1.

* The val function

The val function returns the value of a field described by the model. The val function

20 takes one or more arguments, and the first argument refers to the level-1 data structure

containing the field, and must be either:

* the name of a level-1 data structure in the language model

* the integer 1, indicating the level-1 parent of the variable-size field. In this case, the

variable-size field and the field that specifies its size are in the same data structure, and so

25 have a common level-1 parent.

The subsequent arguments are integers that the specify the ordinal number within its

substructure of the (sub)field that should be dereferenced.

30 By default, COBOL data fields within a structure are not aligned on type-specific

boundaries in storage. For example, the "natural" alignment for a four-byte integer is a

full-word storage boundary. Such alignment can be specified by using the

SYNCHRONIZED clause on the declaration. Otherwise, data fields start immediately

31

after the end of the preceding field in the structure. Since COBOL does not have bit data,
fields always start on a whole byte boundary.

For PL/I, the situation is more complicated. Alignment is controlled by the Aligned and
Unaligned declaration attributes. By contrast with COBOL, most types of data, notably
binary or floating-point numbers, are aligned on their natural boundaries by default.

4.   d.   i)   Formula examples

4.   d.   i)   a)   COBOL

The examples use the proposed inline comment indicator "*>" from the draft standard. It
is not yet legal COBOL usage.

1.   Consider the following data description:

```
*> Field                    Offset

01 Used-Car.                *> "0"
   02 Summary.              *>  "0"
      03 Make pic x(36).     *>    "0"
      03 Model pic x(44).    *>    "36"
      03 VIN pic x(13).      *>    "80"
      03 Color pic x(10).    *>    "93"
         88 Red value 'Red'.
         88 White value 'White'.
         88 Blue value 'Blue'.
   02 History.              *>  "103"
      03 Mileage pic 9(6).   *>   "103"
      03 NumClaims binary pic 9.   *>   "109"
      03 InsCode pic x.       *>   "111"
      03 Claims.             *>   "112"
         04 Claim occurs 1 to 9 times
            depending on NumClaims.  *> stride(1) = "157"
            05 ClaimNo pic x(12).
            05 ClaimAmt binary pic 9(5).
            05 Insurer pic x(38).
            05 Details pic x(100).
   02 Price comp pic 9(5)v99.      *>   "add(112,mpy(val(1,2,2),157))"
```

32

The offset of *Model* is straightforward, and is given by the formula *"36"*. So is that of *Claims*, which is *"112"*.

5   But because the array *Claim* can occur a variable number of times, the structure History is a variable-size field. Thus the offset of *Price*, which immediately follows *Claims*, requires a more complicated formula, involving the array stride (the distance between successive elements along a specific dimension). For *Claim*, there is only one dimension, and the formula for its stride is "154". Thus the formula for the offset of *Price* is:

"add(112,mpy(val(1,2,2),154))"

10

The first argument of the val function is 1, meaning that the field containing the value at run-time, *NumClaims*, is in the same level-1 structure, *Used-Car*, as the field, *Price*, whose offset is specified by the formula. The other two arguments are 2 and 2. The first 2 refers to the second immediate subcomponent, *History*, of *Used-Car*. The second 2

15   means that the field to be dereferenced is the second component of *History*, that is, *NumClaims*.

If the *OCCURS DEPENDING ON* object were in a separate structure, the third subcomponent of level-1 structure *Car-Data*, say, then the val function would be

20   *"val(Car-Data,3)"*.

COBOL structure mapping is top-down, although the direction doesn't make any difference unless the SYNCHRONIZED clause is specified on the data declaration. Specifying SYNCHRONIZED forces alignment of individual fields on their natural

25   boundaries, and thus introduces "gaps" into the structure mapping. Consider the following data structure, which is identical to the previous example, except for the SYNCHRONIZED clause:

```
*> Field                Offset

01 Used-Car SYNCHRONIZED    *> "0"
  02 Summary.               *>  "0"
    03 Make pic x(36).      *>   "0"
    03 Model pic x(44).     *>  "36"
```

```
            03 VIN pic x(13).          *>    "80"
            03 Color pic x(10).        *>    "93"
              88 Red value 'Red'.
              88 White value 'White'.
              88 Blue value 'Blue'.
          02 History.                *>   "103"
            03 Mileage pic 9(6).        *>    "103"
            03 NumClaims binary pic 9.   *>    "110"
            03 InsCode pic x.          *>    "112"
            03 Claims.               *>    "113"
              04 Claim occurs 1 to 9 times
                 depending on NumClaims.   *> stride(1) = "160"
              05 ClaimNo pic x(14).
              05 ClaimAmt binary pic 9(5).
              05 Insurer pic x(39).
              05 Details pic x(100).
          02 Price comp pic 9(5)v99.       *>
          "add(add(113,mpy(val(1,2,2),160)),3)"
```

To position the binary fields on their appropriate half-word or full-word storage boundaries, COBOL introduces padding, known as "slack bytes", into the structure. Working top-down, this padding is introduced immediately before the field needing alignment. So there is one byte of padding between *Mileage* and *NumClaims*.

For an array, such as Claim, COBOL not only adjusts the padding within an element, but also the alignment of each element of the array. In the example, the first occurrence of Claim starts one byte past a full-word boundary. Because the field ClaimNo is an exact number of full-words long, it ends one byte past a full-word boundary, so COBOL inserts three bytes of padding immediately before the binary full-word integer ClaimAmt. And to align subsequent occurrences, so that they too start one byte past a full-word boundary like the first, and can thus have an identical configuration, COBOL adds three bytes of padding at the end of each occurrence.

Finally, after padding, each occurrence of Claim (starts and) ends one byte past a full-word boundary, so COBOL puts three bytes of padding before the binary field Price. As a result of all these extra bytes, the formula for the offset of Price has changed considerably from the unaligned example, and is now:

34

"add(add(113,mpy(val(1,2,2),160)),3)"

There are several differences between the OCCURS DEPENDING ON construct and PL/I's Refer option. Storage for COBOL structures is always allocated at the maximum

5   size, whereas PL/I structures are allocated at the actual size specified by the Refer option. It is legal and usual to change the number of occurrences in a particular instance of a variable-size COBOL array, and this has the effect of changing the location and offset of any fields that follow the array. For PL/I, the value of the Refer object of a particular instance of a structure is intended to be fixed during execution Thus aligned objects

10   following a variable-size field are always correctly aligned for each instance of the structure, because the amount of padding is computed uniquely for each instance, as determined by the Refer option. By contrast, the amount of padding for any aligned fields following a variable-size COBOL array is computed assuming the maximum array size, and is fixed at compile time. If the array is smaller than its maximum size, then the

15   alignment will typically be incorrect. For instance in this example:

> *1 a sync.*
> *2 b binary pic 9.*
> *2 c pic x occurs 1 to 5 times depending on b.*
> *2 d binary pic 9(9).*

20

COBOL inserts one byte between c and d. The alignment of d is therefore correct for only two values of b, the maximum, 5, and 2.

As noted above, the formulas describe not only offsets of fields within a structure, but

25   also properties of arrays, such as bounds and strides. COBOL does not have true multi-dimensional arrays, although element references do use multiple subscripts. Instead, COBOL has arrays of arrays, as in the following simple example:

1 a.	**< offset = "0"*
2 d1 occurs 5 times.	**< offset = "0"*
	*< lbound(1) = "1"
	*< hbound(1) = "5"
	*< stride(1) = "168"
3 d2 occurs 6 times.	**< offset = "0"*

30

SVL920010038US1

< lbound(2) = "1"
< hbound(2) = "6"
< stride(2) = "28"
4 el binary pic 9(9) occurs 7 times. *< offset = "0"*
< lbound(3) = "1"
< hbound(3) = "7"
< stride(3) = "4"

The program can refer to slices of the array by subscripting the higher-level container fields, for example, *d1(2)* or *d2(3, 4)*, but the normal kind of reference is to the low-level elements using the full sequence of subscripts, for instance, *el(4, 5, 6)*. To locate element *el(m, n, o)* using these stride formulas, one would take the address of a and add to it *(m-1)*168 + (n-1)*28 + (o-1)*4*. For COBOL, the lower bound of an array subscript is always 1. That is, the first element is always element(1), and vice versa.

Needless to say, any dimension of the array can have the OCCURS DEPENDING ON clause, and the array can be followed by other fields, which complicates the formulas a lot. Consider the example:

```
1 a.
2 x1 binary pic 9.      *< offset = "0"
2 x2 binary pic 9.      *< offset = "2"
2 x3 binary pic 9.      *< offset = "4"
2 d1 occurs 1 to 5 times  *< offset = "6"
   depending on x1.     *< lbound(1) = "1"
                        *< hbound(1) = "val(1,1)"
                        *< stride(1) = "mpy(val(1,2),mpy(val(1,3),4))"
3 d2 occurs 1 to 6 times *< offset = "6"
   depending on x2.     *< lbound(2) = "1"
                        *< hbound(2) = "val(1,2)"
                        *< stride(2) = "mpy(val(1,3),4)"
4 el binary pic 9(9)    *< offset = "6"
   occurs 1 to 7 times  *< lbound(3) = "1"
   depending on x3.     *< hbound(3) = "val(1,3)"
                        *< stride(3) = "4"
2 b binary pic 9(5).    *< offset = "see below!"
1
```

Computing the address of a particular element still involves the stride formulas, but these are no longer simple integers. The address of element el(m, n, o) in the above example is given by taking the address of a and adding to it:

SVL920010038US1

$(m-1)*stride(1) + (n-1)*stride(2) + (o-1)*stride(3)$, i.e.,
$(m-1)*4*val(1,3)*val(1,2) + (n-1)*4*val(1,3) + (o-1)*4$.

Similarly, these stride formulas are used in the formula for the offset of b:

5   *"add(6,mpy(val(1,1),mpy(val(1,2),mpy(4,val(1,3))))*

4.   d.   i).   b).   PL/I

1.   Given the following structure

```
dcl                        /* offset              */
10    1 c unaligned            /* "0"                */
      ,2 c1             /* "0"               */
        ,3 c2   fixed bin(31)    /*    "0"           */
        ,3 c3   fixed bin(31)    /*    "4"           */
      ,2 c4             /* "8"              */
15      ,3 c5   fixed bin(31)    /*    "0"           */
        ,3 c6   fixed bin(31)    /*    "4"           */
        ,3 c7   fixed bin(31)    /*    "8"           */
      ,2 c8   fixed bin(31)    /* "20"             */
      ,2 c9   char( * refer(c7) ) /* "24"          */
20    ,2 c10  char(6)          /* "add(24,val(1,2,3))"    */
      ,2 c11  char(4)          /* "add(add(24,val(1,2,3)),6)" */
    ;
```

The offset of c3 would be given by the simple formula "4", but the offset of c10 would be

25   given by the formula:

     "add(24,val(1,2,3))"

The first argument in the above val function is 1, which indicates the current structure, c.

30   The subsequent arguments are 2 and 3, indicating that the third element, c7, of the second

level-2 field, c4, is the field to be dereferenced.

The offset of c11 is equal to the offset of c10 plus the length of c10 and would be given

by the following formula:

35   *"add(add(24,val(1,2,3)),6)"*

PL/I structure mapping is not top-down, and this can be illustrated by examining the

mapping of the following structure:

```
dcl                      /* offset               */
  1 a based,               /* "0"                  */
5   2 b,                   /*  "0"                 */
      3 b1 fixed bin(15),    /*   "0"                 */
      3 b2 fixed bin(15),    /*   "2"                 */
      3 b3 fixed bin(31),    /*   "4"                 */
    2 c,                   /*  "add(8,mod(neg(val(1,1,1)),4))" */
10    3 c1 char( n refer(b1)), /*   "0"                 */
      3 c2 fixed bin(31);    /*   "val(1,1,1)"          */
```

The value of b1 is given by val(1,1,1), and in order to put c2 on a 4-byte boundary, PL/I

puts any needed padding before c (yes, not between c1 and c2), and hence the offset of c

15   would be given by the following formula:

$$"add(8,mod(neg(val(1,1,1)),4))"$$

So if b1 contains the value 3, then this formula becomes add(8,mod(neg(3),4)), which

evaluates to 9. I.e., there is one byte of padding between the structure b and the structure

20   c.

The model also uses these formulas to specify the bounds and strides in an array, where

the stride is defined as the distance between two successive elements in an array.

25   For example, in the following structure, the second dimension of a.e has a stride specified

by the formula "4", and the first dimension by the formula "20":

```
dcl
  1 a,                    /* offset = "0"       */
    2 b(4) fixed bin(31),     /* offset = "0"        */
30                            /* lbound(1) = "1"      */
                              /* hbound(1) = "4"      */
                              /* stride(1) = "4"      */
    2 c(4) fixed bin(31),     /* offset = "16"        */
                              /* lbound(1) = "1"      */
35                            /* hbound(1) = "4"      */
                              /* stride(1) = "4"      */
    2 d(4) char(7) varying,   /* offset = "32"        */
                              /* lbound(1) = "1"      */
```

38

```
                        /* hbound(1) = "4"      */
                        /* stride(1) = "9"      */
        2 e(4,5) fixed bin(31);    /* offset = "68"      */
                        /* lbound(1) = "1"      */
                        /* hbound(1) = "4"      */
                        /* stride(1) = "20"     */
                        /* lbound(2) = "1"      */
                        /* hbound(2) = "5"      */
                        /* stride(1) = "4"      */
```

5

10

This means that to locate the element a.e(m,n), one would take the address of a.e and add

to it $(m-1)*20 + (n-1)*4$.

If the example were changed slightly to:

15

```
        dcl
          1 a(4),              /* offset = "0"      */
                        /* lbound(1) = "1"      */
                        /* hbound(1) = "4"      */
                        /* stride(1) = "40"     */
          2 b fixed bin(31),    /* offset = "0"      */
          2 c fixed bin(31),    /* offset = "4"      */
          2 d char(7) varying,  /* offset = "8"      */
          2 e(5) fixed bin(31);  /* offset = "20"     */
                        /* lbound(1) = "1"      */
                        /* hbound(1) = "5"      */
                        /* stride(1) = "4"      */
```

20

25

then there is padding between d and e, but the user of the type descriptor can be blissfully

30   unaware and simply use the stride and offset formulas to locate any given array element.

The stride for a is "40", the stride for e is "4", and the offset for e is "20". This means that

to locate the element a(m).e(n), one would take the address of a and add to it $(m-1)*40 +$

$20 + (n-1)*4$.

35

Finally, if the example were changed again to:

```
        dcl
          1 a(4),              /* offset = "0"      */
                        /* lbound(1) = "1"      */
```

39

```
                                    /* hbound(1) = "4"      */
                                    /* stride(1) = "40"     */
                      2 b fixed bin(31),      /* offset = "0"       */
                      2 c(8) bit(4),          /* offset = "4"        */
5                                   /* lbound(1) = "1"      */
                                    /* hbound(1) = "8"      */
                                    /* stride(1) = "4"      */
                      2 d char(7) varying,    /* offset = "8"        */
                      2 e(5) fixed bin(31);   /* offset = "20"       */
10                                  /* lbound(1) = "1"      */
                                    /* hbound(1) = "5"       */
                                    /* stride(1) = "4"       */
```

then the computations for a.e are the same as above, but the computations for a.c become

15   interesting.

The stride for a is still "40", the stride for c is "4" (but this "4" is a count of bits, not

bytes), and the byte offset for c is "4". To locate the element a(m).c(n), one needs both a

byte address and a bit offset. For the byte address, one would take the address of a and

20   add to it $(m-1)*40 + 4 + ((n-1)*4)/8$. The bit offset of a(m).c(n) would be given by

$mod((n-1)*4,8)$.

4.   e.   Type Descriptor Specification

4.   e.   i.   TDLang Metamodel Specification

25

TDLang Classes - General Overview.   TDLang classes serve as a layer of abstraction

between any CAM language model and the TypeDescriptor model.

30   Since any CAM language model can plug into the TDLang model, the Type Descriptor

model only needs to understand how to interface with TDLang in order to access any

CAM language model.

The TDLang model does not provide any function on its own and therefore only makes

35   sense when it is attached to a language model. TDLang is common to all the CAM

SVL920010038US1

language models and is the type target for the association from TypeDescriptors to the language models.

Note all TDLang classes are abstract and they serve as the base classes to the language

5   metamodels.

TDLangClassifier. TDLangClassifier is the parent class of all language-specific Classifier classes and TDLangComposedType. The TDLangSharedType association is derived from the language's "+sharedType" association from Element to Classifer class.

10   The association should be marked "volatile," "transient," or "derived" to reflect that the association is derived from the language model. The TDLangClassifier is derived from TDLangModelElement

TDLangElement. TDLangElement is the parent class of all language-specific Element

15   classes. The tdLangTypedElement association is derived from the language's "+typedElement" association from Classifer to Element class. The association should be marked "volatile", "transient", and "derived" to reflect that the association is derived from the language model.

20   The tdLangElement association is derived from the language's "+element" association from Classifer to Element class. The association should be marked "volatile," "transient," or "derived" to reflect that the association is derived from the language model.

25   TDLangComposedType. The TDLangComposedType is the parent class of all language-specific ComposedTypes. The TDLangGroup association is derived from the language's "+group" association from Element to ComposedType class. The association should be marked "volatile," "transient," or "derived" to reflect that the association is derived from the language model. The TDLangComposedType is derived from

30   TDLangClassifier.

41

4. e. ii. Type Descriptor Metamodel Specification

The Type Descriptor package defines a model for describing the physical implementation of a data item type. This model is language neutral and can be used to describe the types

5 of many languages. Inspections of the type model for different languages can determine the conformance possibilities for the language types. For example, a long type in Java is often identical to a binary type in COBOL, and if so, the types may be interconverted without side effect. On the other hand, an alphanumeric type in COBOL is fixed in size and if mapped to a Java type, will lose this property. When converted back from Java to

10 COBOL, the COBOL truncation rules may not apply, resulting in computation anomalies.

AggregateInstanceTD. For each instance of an aggregate, there is an instance of this class. To find the children of this aggregate, one must navigate the associations back to

15 language Classifier then downcast to language Composed Type and follow the association to its children.

  Derived from InstanceTDBase

  Public Attributes:

   union : boolean = false

20 Distinguishes whether the aggregation is inclusive (e.g. a structure) or exclusive (e.g. a union).

If union = true, storage might be overlaid and as a result the interpretation of the content may be uncertain.

25

ArrayTD. ArrayTD holds information for array types.

  Public Attributes:

   arraryAlign : int

  Alignment requirements in addressible units for each element in the array.

30 strideFormula : string

42

A formula that can be used to calculate the displacement address of any element in the array, given an index.

strideInBit : boolean

5 True indicates strideFormula value in bits

False indicates strideFormula value in bytes

upperBoundFormula : String

Declared as a String for instances when this value is referenced by a variable.

This attribute supplies the upperbound value of a variable size array

10 Upperbound is required when traversing back up the entire structure.

lowerBoundFormula : String

Declared as a String for instances when this value is referenced by a variable.

This attribute supplies the lowerbound value of a variable size array.

15 InstanceTDBase. InstanceTD has instances for each declared variable and structure element.

To find the parent of any instance (if it has one) one must navigate the associations back to TDLangElement, follow the association to TDLangClassifier to locate the parent, then follow the associations back into the TypeDescriptor model.

20 Public Attributes:

offsetFormula : string

A formula for calculating the offset to the start of this instance.

This attribute is String because this field may not always be an integer value. For

25 example, (value(n) + 4) could be a possible value.

NOTE: The offset value is calculated from the top-most parent. (e.g., for a binary tree A-->B, A-->C, B-->D, B-->E. The offset to D is calculated from A to D, not B to D)

30 contentSizeFormula : string

Formula for calculating the current size of the contents

43

allocSizeFormula : string

Formula for calculating the allocated size of the instance

formulaInBit : boolean

True indicates offsetFormula, contentSizeFormula, and allocSizeFormula values

are in bits

False indicates offsetFormula, contentSizeFormula, and allocSizeFormula values

are in bytes

defaultEncoding : String

Physical encoding - how many bits used to encode code points, how are the code

points mapped onto bit patterns

Contains info on how string content is encoded: EBCDIC, ASCII, UNICODE,

UTF-8, UTF-16, codepage numbers, etc...

accessor : enumeration

Specifies access-rights for this element.

defaultBigEndian : boolean

True if this element is Big Endian format.

floatType : enumeration

Specifies this element's float type.

PlatformCompilerType.   A specific data type for a particular platform and compiler.

NOTE: There needs to be some way to identify the platform and compiler. This class can

be specializedor have an attribute, or be simplified by putting an attribute on

InstanceTDBase.

Public Attributes:

platformCompilerType : String

This attribute specifies the type of compiler used to create the data in the

language model. Usage of this string is as follows:

"Vendor name, language, OS, hardware platform" (e.g., "IBM, COBOL, OS390,

390" or "IBM, PLI, WinNT, Intel")

44

SimpleInstanceTD. An instance of a Simple type in the language model.
Derived from InstanceTDBase

NumberTD

5   All numbers representations.
           Currently includes Integers and Packed Decimals

           Note required fields for these types of Numbers:

10   *Integer*
           Base=2
           MSBU=0 or 1
           Signed/Unsigned
           Size (in bytes) = 1,2,4,8 (16)

15
           Packed Decimal
           Base=10
           MSBU=0
           Signed

20   Width=1-63

           Float
           Base=2(IEEE), 16(Hex)
           MSBU=0 or 1

25   Signed
           Size (in bytes)=4,8,16
           Encoding Details...

           Derived from BaseTD

30
           Public Attributes:

45

base : int

The base representation of this number.   2=binary, 10=decimal, 16=hex,

...

baseWidth : int

Number of bits used to represent base:

e.g. binary = 1, decimal = 8, packed = 4

baseInAddr : int

Number of baseWidth units in addressable storage units - e.g. decimal = 1, packed = 2, binary = 8 where the addressable unit is a byte. If the addressable unit was a 32 bit word, decimal would be 4, packed would be 8, and binary would be 32.

baseUnits : int

Number of base units in the number. This times the base width must be less than or equal to the width times the addrUnit.

For example, a 24 bit address right aligned in a 32 bit word would have base = 1, basewidth=24, baseInAddr=8, width=4.

signCoding : enumeration

A set of enumerations - 2's complement, 1's complement, and sign magnitude for binary; zone signs for decimal, packed signs for packed decimal, and unsigned binary, unsigned decimal.

checkValidity : boolean

True if language model is required for picture string support

packedDecimalSign : enumeration

Used to determine the code point of the sign in COBOL decimal numbers, where the sign is combined with the leading or trailing numeric digit.

FloatTD

Floating points

Derived from BaseTD

Public Attributes:

floatType : enumeration

Specifies this element's float type.

1 StringTD

5   Alphanumeric type

Derived from BaseTD

Public Attributes:

encoding : String

10   Physical encoding - how many bits used to encode code points, how are

the code points mapped onto bit patterns

Contains info on how string content is encoded: EBCDIC, ASCII,

UNICODE, UTF-8, UTF-16, codepage numbers, etc...

15   lengthEncoding : enumeration

Possible values for lengthEncoding:

- Fixed length (where total length equals declared string length)

- Length prefixed (where total length equals declared string length plus

length of header bytes; usually 1, 2, or 4 bytes)

20   - Null terminated (known as varyingZ strings in PL/I) (where a null

symbol is added to the end of string; so the maximum length could be up

to declared string length plus one byte to represent null character)

maxLengthFormula : String

Formula specifying the maximum length of this string.

25   checkValidity : boolean

True if language model is required for picture string support

textType : String = Implicit

Value is 'Implicit' or 'Visual'

orientation : StringTD = LTR

30   Value is 'LTR', 'RTL', 'Contextual LTR', or 'Contextual RTL'

where LTR = Left to Right

47

and RTL = Right to Left

Symmetric : boolean = true

True if symmetric swapping is allowed

numeralShapes : String = Nominal

5   Value is 'Nominal', 'National', or 'Contextual'

textShape : String = Nominal

Value is 'Nominal', 'Shaped', 'Initial', 'Middle', 'Final', or 'Isolated'

AddressTD

10   Type to represent pointers/addresses

Rationale for this class:

Addresses should be considered separate from NumberTD class because some languages

on certain machines (e.g., IBM 400) represent addresses with additional information,

such as permission type (which is not represented in NumberTD class)

15   Derived from BaseTD

Public Attributes:

permission : String

Specifies the permission for this address. Used primarily for AS/400 systems.

20   bitModePad : enumeration

Specifies the number of bits for this address. Used to calculate padding.

BaseTD

The base class of typeDescriptor.

25

The BaseTD model type is specialized to hold implementation

information which is common for all data elements of the same language

type. The information which describes a 32 bit signed binary integer on a

specific hardware/software platform is thus instantiated only once in a

30   given model instantiation, no matter how many data elements may be

declared with this type.

Public Attributes:

addrUnit : enumeration

Number of bits in storage addressable unit

  - bit/byte/word/dword

5

width : int

Number of addressable storage units in the described type. This can be 1,

8, 16, 32 bits.

alignment : int

10   Required alignment of type in address space - e.g. word aligned 32 bit

integer would have alignment of 4

nickname : int

Name of the base element. This should uniquely identify an instance of a

simple type to allow logic based on name rather than logic based on

15   combinations of attributes. E.g. "S390Binary31_0" for a 32 bit S/390

unscaled binary integer

bigEndian : boolean

True if this element is Big Endian format.

20   Stereotypes

bitModePad

Public Attributes:

16bit :

25   24bit :

31bit :

32bit :

64bit :

128bit :

30

signCoding

Note that this model does not include the following COBOL usages:

1   POINTER

PROCEDURE-POINTER

OBJECT REFERENCE

5   Public Attributes:

twosComplement :

onesComplement :

signMagnitude :

zoneSigns :

10   packedSigns :

unsignedBinary :

unsignedDecimal :

lengthEncoding

15   Public Attributes:

fixedLength :

lengthPrefixed :

nullTerminated :

20   floatType

Public Attributes:

Unspecified :

IEEEextendedIntel :

For IEEE extended floats running on Intel Platform

25   IEEEextendedAIX :

For IEEE extended floats running on AIX Platform

IEEEextendedOS390 :

For IEEE extended floats running on OS/390 Platform

IEEEextendedAS400 :

30   For IEEE extended floats running on AS400 Platform

IEEEnonextended :

50

For IEEE non-extended floats

IBM390Hex :

For Hexadecimal floats running on IBM OS/390

IBM400Hex :

5   For Hexadecimal floats running on IBM AS400

accessor

Public Attributes:

ReadOnly :

10   WriteOnly :

ReadWrite :

NoAccess :

packedDecimalSign

15   Public Attributes:

MVS :

MVSCustom :

NT/OS2/AIX :

20

5.   Illustrative Applications of the Common Application Metamodel and System

Various complex transaction, occurring across a plurality of individual platforms require

the seamlessness and transparency of the Common Application Metamodel. These

25   transactions include "rich" transactions, high level group ware, and financial transactions.

FIGURE 13 illustrates a simplified network configuration for a "rich" transaction where,

for example, an order is entered at a terminal, and is passed to and through a Web server

to a manufacturer's application server. The manufacturer's application server searches

30   through it's own database server, as well as its vendors' dissimilar and incompatible

database servers and databases, transparently connected by the connectors described

51

herein, querying for statuses, prices, and delivery dates, of components, and placing orders for needed components to satisfy the order.

5   Strictly for purposes of illustration and not limitation, a rich transaction is illustrated with the purchase of an automobile, from configuring and ordering the new automobile through financing the automobile, collecting the elements and components the new automobile, assembling the automobile, and delivering it to the customer. In configuring the automobile, consider the inclusion of, e.g., a traction control module, which may require the inclusion of one sub-set of engines, while precluding the inclusion of another

10   sub-set of engines, for example, for technical or performance reasons. Similarly, the selection of one color may preclude the selection of certain upholstery combinations, for example, for reasons of inventory or availability. Of course, if you don't select an engine and a transmission, you don't have an automobile. Certain choices must be made. The problem is one of "If you pick `X`, you must also pick `Y` and `Z`, but if you pick `Y`,

15   you can not get `A` or `B`, and you have already selected `A`." That is, selection of one component or sub-system may remove some previously selected components or sub-systems from the system. After configuring the car, frequently based on the availability of selected elements in the supply pipeline, the order is sent to the manufacturer's server for fulfillment, including manufacturing scheduling. The manufacturer would query the

20   vendors in the supply chain, for example, a transmission vendor, who would, in turn, query controller vendors, gear vendors, housing vendors, hydraulic conduit vendors, and the like. These vendors would, in turn query their sub-vendors.

Using the connector method and system, the dealer and customer would configure and

25   order the car at a terminal 3901. The configuration and order entry would be sent to the manufacturer's server 3911. The manufacturer's server queries the servers 3921, 3931, and 3441, (where server 3921 could be the server for the manufacturer's in house components business and for its purchasing function) or its direct vendors, who in turn query third sub-vendors, 3923, 3933, and 3933. The queries could be in the form of "fill

30   in the blank" SQL queries, subject to access authorizations. The sub-vendor servers, 3923, 3933, and 3443 specialized views to the vendor's servers 3921, 3931, 3441, as

SVL920010038US1

clients, which in turn would present specialized views to the manufacturer's server 3911 as a client. The data returned to the manufacturer's server 3911 could be optimized with optimization and production scheduling applications, to present a serial number and delivery date to the buyer at the dealer's server.

5

A further application of the connectors of the present invention is in connection with groupware, and especially linking "islands of groupware" together. "Groupware" is a broad term applied to technology that is designed to facilitate the work of groups. Groupware technology may be used to communicate, cooperate, coordinate, solve problems, compete, or negotiate. A Groupware suite of applications is comprised of programs that help people work together collectively, even if located remotely from each other. Groupware services can include the sharing of calendars, collective writing, e-mail handling, shared database access, electronic meetings, video conferencing, and other activities.

15

FIGURE 14 illustrates a group ware session spread across multiple group ware applications running on multiple, disparate platforms, and connected by the connectors described herein. Specifically illustrated are two groupware systems 4010 and 4020. Each system contains e-mail enables applications, 4013 and 4023, such as e-mail, schedule/calendar, word processor, spread sheet, graphics, and CTI (computer telephony integration). These are supported by message API's 4015, 4025 and operating systems, 4017 and 4027, and local groupware servers 4011 and 4021. Each of the local groupware servers, 4011 and 4021, has a connector of the present invention, an e-mail server, an e-mail data base, a directory server, and a replication database. The two groupware servers 4011 and 4021 communicate over a telecommunications medium, as the Net, a LAN, a WAN, or even the PSTN.

E-mail is by far the most common and basic groupware application. While the basic technology is designed to pass simple messages between people, even relatively simple e-mail systems today typically include features for forwarding messages, filing messages, creating mailing groups, and attaching files with a message. Other features that have been

explored include content based sorting and processing of messages, content based routing of messages, and structured communication (messages requiring certain information).

Workflow systems are another element of groupware. Work flow systems allow

5   documents to be routed through organizations through a relatively-fixed process. A simple example of a workflow application is an expense report in an organization: an employee enters an expense report and submits it, a copy is archived then routed to the employee's manager for approval, the manager receives the document, electronically approves it and sends it on and the expense is registered to the group's account and

10   forwarded to the accounting department for payment. Workflow systems may provide features such as routing, development of forms, and support for differing roles and privileges.

Hypertext is a system for linking text documents to each other, with the Web being an

15   obvious example. However, whenever multiple people author and link documents, the system becomes group work, constantly evolving and responding to others' work. Another common multi-user feature in hypertext (that is not found on the Web) is allowing any user to create links from any page, so that others can be informed when there are relevant links that the original author was unaware of.

20

Group calendaring is another aspect of groupware and facilitates scheduling, project management, and coordination among many people, and may provide support for scheduling equipment as well. Typical features detect when schedules conflict or find meeting times that will work for everyone. Group calendars also help to locate people.

25

Collaborative writing systems may provide both realtime support and non-realtime support. Word processors may provide asynchronous support by showing authorship and by allowing users to track changes and make annotations to documents. Authors collaborating on a document may also be given tools to help plan and coordinate the

30   authoring process, such as methods for locking parts of the document or linking separately-authored documents. Synchronous support allows authors to see each

54

other's changes as they make them, and usually needs to provide an additional communication channel to the authors as they work.

Group may be synchronous or real time, such as shared whiteboards that allow two or
5   more people to view and draw on a shared drawing surface even from different locations. This can be used, for instance, during a phone call, where each person can jot down notes (e.g. a name, phone number, or map) or to work collaboratively on a visual problem. Most shared whiteboards are designed for informal conversation, but they may also serve structured communications or more sophisticated drawing tasks, such as collaborative
10   graphic design, publishing, or engineering applications. Shared whiteboards can indicate where each person is drawing or pointing by showing telepointers, which are color-coded or labeled to identify each person.

A further aspect of groupware is video conferencing. Video conferencing systems allow
15   two-way or multi-way calling with live video, essentially a telephone system with an additional visual component, with time stamping for coordination.

Decision support systems are another aspect of groupware and are designed to facilitate groups in decision-making. They provide tools for brainstorming, critiquing ideas,
20   putting weights and probabilities on events and alternatives, and voting. Such systems enable presumably more rational and evenhanded decisions. Primarily designed to facilitate meetings, they encourage equal participation by, for instance, providing anonymity or enforcing turn-taking.

25   Multi-player games have always been reasonably common in arcades, but are becoming quite common on the Internet. Many of the earliest electronic arcade games were multi-user, for example, Pong, Space Wars, and car racing games. Games are the prototypical example of multi-user situations "non-cooperative", though even competitive games require players to cooperate in following the rules of the game. Games can be enhanced
30   by other communication media, such as chat or video systems.

SVL920010038US1

Some product examples of groupware include Lotus Notes and Microsoft Exchange, both of which facilitate calendar sharing, e-mail handling, and the replication of files across a distributed system so that all users can view the same information.

5   What makes the connectors of the present invention particularly attractive for groupware applications is the diversity of groupware offerings and group server architectures, implementations, and platforms. The connector of the invention can act as a front end or gateway to the groupware server.

10   FIGURE 15 illustrates a commercial transaction where real goods are shipped from a seller to a buyer, and various forms of electronic payment and secured electronic payment are used by the buyer to pay the seller, with banks and financial institutions connected through the connectors described herein. Specifically, a merchant or manufacturer 4101 sells a product to a customer 4103 that he has no "history" with. The product is shipped

15   4605. However, the buyer 4103 does not wish to be parted from his money until the goods are received, inspected, approved, and "accepted", while the seller 4101 does not want to give up control of the goods until he has been paid. This fundamental commercial conflict has led to various paradigms, most involving hard copy "near moneys" or instruments of one form or another. Today, the financial transactions are most frequently

20   those involving electronic fund transfers, and electronic versions of notes, instruments, negotiable instruments, documentary drafts, payment orders, letters of credit, warehouse receipts, delivery orders, bills of lading, including claims on goods, that is, documents of title that purport to transfer title or physical custody of goods to the bearer or to a named person, and security interests in goods. Typically, the customer 4103 executes an

25   instrument in favor of the seller 4101, directing the buyer's bank 4121 to pay buyer's money to the seller 4101 through seller's bank 4125. Normally, this is a simple electronic transaction between buyers and sellers who have dealt with each other before, dealing through banks or financial intermediaries who have dealt with each other before, and who are using the same or compatible software applications. However, in the

30   extraordinary case where these preconditions are not satisfied, the connectors of the invention facilitate the electronic, bank-to-bank, side of the transaction.

While the invention has been described and illustrated with respect to applications having a single level of application program interfaces or converters, it is, of course, to be understood that such application program interfaces or converters may be present at

5   various levels in the processing hierarchy, for example between Web Clients and Web servers, between web servers and application servers, between application servers and database servers, and between application servers or database servers or both and various specialized repositories.

10 It is also to be understood, that while the invention has been described and illustrated with respect to certain preferred embodiments and exemplification having individual clients and individual servers, there may be multiple clients, multiple servers, and applications that function as both clients and servers, as exemplified by groupware applications, and there might be multiple parallel lines and/or multiple hierarchical levels

15 of application servers, data servers, and databases, as in systems for rich transactions.

While the invention has been described with respect to certain preferred embodiments and exemplifications, it is not intended to limit the scope of the invention thereby, but solely by the claims appended hereto.

SVL920010038US1